

A data-driven method for improving the correlation estimation in serial ensemble Kalman filters

Michèle De La Chevrotière *

Department of Mathematics, the Pennsylvania State University, USA

John Harlim

*Department of Mathematics and Department of Meteorology and Atmospheric Science, the
Pennsylvania State University, USA*

**Corresponding author address:* Department of Mathematics, the Pennsylvania State University,
109 McAllister Building, University Park, PA 16802-6400, USA

E-mail: mdelachev@psu.edu

ABSTRACT

A data-driven method for improving the correlation estimation in serial ensemble Kalman filters is introduced. The method finds a linear map that transforms, at each assimilation cycle, the poorly estimated sample correlation into an improved correlation. This map is obtained from an offline training procedure without any tuning as the solution of a linear regression problem that uses appropriate sample correlation statistics obtained from historical data assimilation outputs. In an idealized OSSE with the Lorenz-96 model and for a range of linear and nonlinear observation models, the proposed scheme improves the filter estimates, especially when the ensemble size is small relative to the dimension of the state space.

1. Introduction

It is a well known fact (Evensen 1994) that when small ensemble sizes are used in an ensemble Kalman filter (EnKF), the estimated correlation tends to overestimate the correlations between two variables that are far away even if the true correlation is small (Lorenc 2003). The typical approach to address this issue is by zeroing (and damping) the correlation coefficients of longer (intermediate) distances, which practically means ignoring (weighting less) such observations. Such practice is known as *localization* (Houtekamer and Mitchell 1998). Many localization functions have been proposed, with tapering (Furrer and Bengtsson 2007), the Gaspari-Cohn function (Gaspari and Cohn 1999), or other distance-based localization functions for multiphase flows (Chen and Oliver 2010). While these approaches have been very useful, they require exhaustive tuning which can be expensive for problems with multiscale spatial correlations (Dong et al. 2011; Kirchgessner et al. 2014; Flowerdew 2015). Furthermore, since the correlation between a state variable x and a nonlinearly observed state $h(x)$ may not necessarily be a function of spatial distance, specifying a localization function for such correlations may not be trivial. Alternatively, a class of adaptive (or “flow-dependent”) methods has also been proposed, however they depend on many parameters such as the integer powers in Bishop and Hodyss (2007), or require specification of a prior distribution for the sampling correlation (e.g. see Zhen and Zhang 2014). There is also a nonparametric method (Jun et al. 2011) which provides a guideline for estimating the background covariance matrix.

More recently, an approach for estimating the localization function that requires almost no tuning has been proposed by Anderson and Lei (2013). This approach consists of training empirical localization functions (ELF) from an observation system simulation experiment (OSSE) data assimilation output by minimizing the analysis error, i.e., the difference between the analysis mean

update and the truth (or observations when the truth is unknown). Numerically, they obtain ELF in stages: The first ELF is obtained by minimizing the analysis error from an EnKF without localization, the second ELF is obtained by minimizing the analysis error from EnKF using the first ELF estimate, and so on. In this sense, there is still some tuning required, for instance in determining the stopping criteria.

Motivated by this approach, we introduce a data-driven method for improving the sample correlation estimation in the serial EnKF when small ensemble sizes are used. In particular, we consider finding a linear map that takes poorly estimated sample correlations from each EnKF cycle and transforms them into an improved correlation matrix. This linear map is computed by regressing appropriate sample correlation statistics from historical data assimilation output obtained in a training phase and we will show that there is no tuning required. We shall see that this linear map generalizes the standard localization function, and our method for specifying this map can be interpreted as a discrete (or Monte-Carlo) approximation to a minimization problem with cost function inspired by Anderson (2012) and Flowerdew (2015). Our approach is relatively easy to implement in any observation network and it does not parameterize the prior distribution as in Zhen and Zhang (2014) or the likelihood function as proposed in Anderson (2016).

The remainder of this paper is organized as follows: In Section 2, we review the basic statistical results on sampling error in estimating correlation which motivate the current approach. In Section 3, we introduce our new algorithm for specifying the linear map and describe its implementation with a serial EnKF of Anderson (2003). In Section 4, we demonstrate the numerical results in an idealized OSSE setup using the Lorenz-96 model (Lorenz 1996), and compare it to an EnKF with the well-tuned Gaspari-Cohn localization function as well as ETKF with a very large ensemble size, where we use the ETKF formulation of Hunt et al. (2007) without localization. We close this paper with a short summary in Section 5.

2. Sampling error in correlation statistics of EnKF

The goal of this section is to give readers a brief review to understand sampling errors in an EnKF. In particular, we will show that it is rather difficult to provide a practically useful estimate for sampling errors in the EnKF in the linear and Gaussian case, and discuss how this issue is compounded in high-dimensional nonlinear filtering problems. While the detailed formulation in this discussion may seem to be a digression, it aims to provide intuition for the algorithm proposed in Section 3.

Consider the discrete-time linear and Gaussian filtering problem,

$$\mathbf{x}_{m+1} = F\mathbf{x}_m + \boldsymbol{\zeta}_{m+1}, \quad \boldsymbol{\zeta}_m \sim \mathcal{N}(0, Q), \quad (1)$$

$$\mathbf{y}_m^o = H\mathbf{x}_m + \boldsymbol{\epsilon}_m, \quad \boldsymbol{\epsilon}_m \sim \mathcal{N}(0, R), \quad (2)$$

where F denotes the deterministic model operator and Q denotes the Gaussian system noise covariance matrix. The time index is defined as $t_{m+1} = t_m + \tau$, where τ denotes the observation time step. The optimal solution (in the least-square sense) to this filtering problem, assuming both noises are uncorrelated, $\mathbb{E}[\boldsymbol{\epsilon}_m \boldsymbol{\zeta}_n^\top] = 0$ for all m, n , is given by the Kalman filter formula,

$$\begin{aligned} \bar{\mathbf{x}}_m^a &= \bar{\mathbf{x}}_m^f + \mathcal{K}_m(\mathbf{y}_m^o - H\bar{\mathbf{x}}_m^f), & P_m^a &= (I - \mathcal{K}_m H)P_m^f, \\ \mathcal{K}_m &= P_m^f H^\top (H P_m^f H^\top + R)^{-1} \\ \bar{\mathbf{x}}_{m+1}^f &= F\bar{\mathbf{x}}_m^a, & P_{m+1}^f &= F P_m^a F^\top + Q \end{aligned} \quad (3)$$

where $\bar{\mathbf{x}}_m^f, P_m^f$ denote the forecast mean and covariance estimates before accounting for the observations \mathbf{y}_m^o , and $\bar{\mathbf{x}}_m^a, P_m^a$ denote the analysis mean and covariance estimates after accounting for the observations \mathbf{y}_m^o . In this rather restrictive scenario, the time asymptotic limit of the Kalman filter covariance solutions, P_m^a , is the solution P^* of the discrete Riccati equation

$$P^* = (I - \mathcal{K}H)(F P^* F^\top + Q),$$

where $\mathcal{K} = P^*H^\top(HP^*H^\top + R)^{-1}$ denotes the limiting Kalman gain, i.e. we have $P_m^a \rightarrow P^*$ as $m \rightarrow \infty$.

The main idea behind ensemble Kalman filters (e.g. Evensen 1994; Anderson 2001; Bishop et al. 2001) is to approximate the forecast distribution by an ensemble of solutions, use the Kalman filter formula in (3) to transform the first two empirical prior moments, $\hat{\bar{x}}_m^f, \hat{P}_m^f$ (obtained from the ensemble average) to analysis empirical mean and covariance, $\hat{\bar{x}}_m^a, \hat{P}_m^a$, and finally draw a Gaussian sample $\{x_m^{a,k}\}$ such that

$$\hat{\bar{x}}_m^a = \frac{1}{K} \sum_{k=1}^K x_m^{a,k} \quad \text{and} \quad \hat{P}_m^a = \frac{1}{K-1} X_m^a X_m^{a\top},$$

where each column of the matrix X_m^a consists of the perturbation of the k^{th} ensemble member from its sample mean, $X_m^{a,k} = x_m^{a,k} - \hat{\bar{x}}_m^a$, and K denotes the ensemble size. Here, m is the time index.

In the Gaussian and linear case, the EnKF solutions are optimal when the sample mean estimate, $\hat{\bar{x}}_m^a$, matches the Kalman mean estimate \bar{x}_m^a and the sample covariance estimate, \hat{P}_m^a , matches the asymptotic covariance P^* as $m \rightarrow \infty$. In such a case, the EnKF analysis ensemble perturbations, $X_m^{a,k}$, are samples of a time independent Gaussian distribution with mean zero and covariance matrix P^* or correlation matrix $\rho = D^{-1/2}P^*D^{-1/2}$, where D is a diagonal matrix whose diagonal components are the diagonal components of P^* . That is, as $m \rightarrow \infty$, the EnKF numerically estimates P^* using the sample covariance matrix \hat{P}_m^a at each data assimilation step m .

From basic statistical theory (Casella and Berger 2002), we know that the standard deviation of the estimated variance \hat{P}_{ii}^a obtained from K independent samples of a scalar Gaussian random variable x_i , the i^{th} component of x , is $\sigma_i = \sqrt{\frac{2}{K-1}P_{ii}^*}$, which is smaller than the variance P_{ii}^* when the ensemble size $K > 3$. In operational implementations, where the ensemble size $K \approx 10 - 100$, this standard deviation is quite small. Therefore, it is reasonable to assume that the sample

covariance \hat{P}_{ii}^a is an accurate estimation of the true variance P_{ii}^* . A similar assumption is taken in Anderson (2012, 2016).

Now let's turn our attention to the sampling errors in the estimation of the non-diagonal components of P^* or essentially the correlation matrix ρ , assuming that the diagonal component of P^* can be accurately estimated by the EnKF solutions $\{x_m^{a,k}\}_{k=1,\dots,K}$. In other words, we are estimating ρ with a sample correlation matrix $r_m^K \equiv D^{-1/2} \hat{P}_m^a D^{-1/2}$, assuming that $D = \text{diag}(P^*) = \text{diag}(\hat{P}_m^a)$, using the EnKF with an ensemble size K . Of course these samples are not temporally independent for observation time interval that satisfies $\tau < T_c$, where T_c denotes the decaying time scale of the underlying dynamics. For $m \rightarrow \infty$ and large enough $\tau > T_c$ such that $X_m^{a,k}$ are temporally independent samples of a multivariate Gaussian distribution with mean zero and covariance P^* , the distribution and moments of its sample correlation coefficient r_m^K are explicitly known (Hotelling 1953). In particular, $\mathbb{E}[(r_{ij}^K - \rho_{ij})^2] = (1 - \rho_{ij}^2)^2 K^{-1} + \mathcal{O}(K^{-2})$, which means that the error is on the order of K^{-1} when $\rho_{ij} = 1$ and $K^{-1/2}$ otherwise. Of course, this error estimate is valid only when $x^{a,k}$ are independent samples of a Gaussian prior and the filtering problem is linear, and so it is not useful in practice.

In general high-dimensional non-Gaussian and nonlinear settings, the problem becomes much harder since the true correlation ρ_m is time-dependent (it does not equilibrate to ρ). Furthermore, the true correlations ρ_m are statistics with respect to the true filtering distribution $p(x_m | y_m^o)$ which is almost always non-Gaussian in the general nonlinear filtering problem. The bottom line is that if one approximates the nonlinear filtering problem with an EnKF, the sample correlation is not the same as the true correlation even in the limit of large ensemble size,

$$\lim_{K \rightarrow \infty} r_m^K \equiv r_m \neq \rho_m, \quad \forall m,$$

where r_m^K denotes the sample correlation obtained from an EnKF with ensemble size K , r_m denotes the limiting correlation as $K \rightarrow \infty$, and ρ_m denotes the true correlation. While estimating the true non-Gaussian statistics ρ_m may be very difficult in practice, it is not hopeless to estimate the limiting sample correlation r_m from a sample correlation r_m^K that is obtained from an EnKF with a finite ensemble size K . This is precisely what we will discuss in the next section, that is, we propose to find a linear map \mathcal{L} that takes the sample correlation r_m^K to the limiting sample correlation r_m . In a special case, this map is exactly what is commonly known as the *localization function*.

3. Algorithm

The main idea in Anderson (2012) and Flowerdew (2015) is to find a localization function by minimizing the mean square difference of the analysis from the truth. Motivated by their ideas and the theoretical understanding above, we propose to find a linear map $\mathcal{L} \in \mathbb{R}^{N \times N \times M}$, a three-dimensional array which transforms the sample correlations between \mathbf{x} and \mathbf{y} , to the corresponding limiting sample correlation. We will learn this map from historical data assimilation outputs as discussed in Section 3a below. Later on, in Section 3b, we show a straightforward modification of a serial EnKF (Anderson 2003) using this linear map \mathcal{L} .

a. Learning the “localization” map

First, let us set our notation. To simplify the notation in the remainder of this paper, we suppress the subscript time m and only use it when necessary. Since we are interested in filtering problems with general observation functions $h : \mathbb{R}^N \rightarrow \mathbb{R}^M$, we consider a sample cross correlation between \mathbf{x} and \mathbf{y} as follows,

$$r_{\mathbf{x}\mathbf{y}}^K = D_{X^a}^{-1/2} \left(\frac{1}{K-1} X^a Y^{a\top} \right) D_{Y^a}^{-1/2} \in \mathbb{R}^{N \times M}, \quad (4)$$

where X^a is defined as before, and the columns of Y^a are defined as $\mathbf{Y}^{a,k} = h(\mathbf{x}^{a,k}) - \hat{\mathbf{y}}^a$, where $\hat{\mathbf{y}}^a = \sum_{k=1}^K h(\mathbf{x}^{a,k})/K$. In (4), the matrices $D_{X^a} = \text{diag}(X^a X^{a\top})/(K-1)$ and $D_{Y^a} = \text{diag}(Y^a Y^{a\top})/(K-1)$ are diagonal. Notice that in the case of linear observation operators $h(\mathbf{x}) = H\mathbf{x}$, we have $Y^a = HX^a$. In the remainder of this article, we suppress the subscript $\mathbf{x}\mathbf{y}$ and understand that $r^K = r_{\mathbf{x}\mathbf{y}}^K$ with limiting correlation $r = \lim_{K \rightarrow \infty} r^K$. We also use the notation $\mathcal{L}(\cdot, i, j) \in \mathbb{R}^N$ to denote an N -dimensional column vector corresponding to fixing the second and third component in the array \mathcal{L} by i and j , respectively.

We consider finding a linear functional $\mathcal{L}(\cdot, i, j) : [-1, 1]^N \rightarrow [-1, 1]$ that takes the poorly sampled correlations between all state variables and the observation j , and maps them to an improved correlation $\tilde{r}(i, j)$ for the pair of state variable i and observation j , defined as $\tilde{r}(i, j) = r^K(\cdot, j)^\top \mathcal{L}(\cdot, i, j)$. We propose to find this map by solving

$$\min_{\mathcal{L}(\cdot, i, j)} \int_{[-1, 1] \times [-1, 1]} (r^K(\cdot, j)^\top \mathcal{L}(\cdot, i, j) - r(i, j))^2 p(r^K|r) p(r) dr^K dr, \quad (5)$$

such that $\tilde{r}(i, j)$ best approximates $r(i, j)$, the limiting correlation between the state variable i and the observation j . Here $p(r)$ is the distribution of the limiting sample correlation whereas $p(r^K|r)$ is the conditional distribution of the sample correlation r^K given r , where we assume that these distributions are stationary. In this formulation, the goal is to find an estimator for the limiting correlation $r(i, j)$ through a linear map which combines information from all the sample correlations $r^K(\cdot, j)$, whereas the approach in (Anderson 2012; Flowerdew 2015) considers a scalar factor applied to the corresponding single sample correlation $r^K(i, j)$ only. Their approach is similar to assuming that the map $\mathcal{L}(\cdot, \cdot, j) \in \mathbb{R}^{N \times N}$ is diagonal in our formulation (see the discussion at the end of this section).

To realize Eqn. (5) numerically, we consider the following training procedure. Consider an OSSE historical data assimilation analysis ensemble, $\{\mathbf{x}_m^{a,k}\}$, where the index $m = 1, \dots, T$ denotes

time and the index $k = 1, \dots, L$ denotes ensemble member, obtained from a large ensemble of size $L \gg 1$, such that $r^L \approx r$. Here, T denotes the number of data assimilation cycles in the training period. Given this analysis ensemble, one can sample the limiting correlation by r_m^L , and also obtain subsampled correlations $\{r_{m,s}^K\}$, using only K members out of L . Here we use the index $s = 1, \dots, S$ to denote the number of possible ways to choose K members out of L , thus one can sample as many as $\max(S) = \binom{L}{K}$. Therefore, from the historical analysis ensemble, we obtain samples $r_m^L \sim p(r)$ and $r_{m,s}^K \sim p(r^K | r = r_m^L)$. We can then apply a Monte-Carlo approximation to the integral in (5) such that the minimization problem becomes

$$\min_{\mathcal{L}(\cdot, i, j)} \frac{1}{TS} \sum_{m,s=1}^{T,S} (r_{m,s}^K(\cdot, j)^\top \mathcal{L}(\cdot, i, j) - r_m^L(i, j))^2, \quad (6)$$

which is essentially a least square problem. To see this, define $A \in \mathbb{R}^{TS \times N}$, where each row of A is $r_{m,s}^K(\cdot, j)^\top$, $\mathbf{u} = \mathcal{L}(\cdot, i, j) \in \mathbb{R}^N$, and the vector $\mathbf{b} \in \mathbb{R}^{TS}$ whose components are $r_{m,s}^L(i, j)$, where $r_{m,s}^L(i, j) = r_m^L(i, j)$ for any m, s . Then the linear regression problem in (6) can be rewritten in compact form as $\min_{\mathbf{u}} (\mathbf{A}\mathbf{u} - \mathbf{b})^\top (\mathbf{A}\mathbf{u} - \mathbf{b})$, ignoring the scalar constant factor $(TS)^{-1}$. An explicit solution to this linear regression problem is given by $\mathbf{u} = (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{b}$. Repeating this procedure for every indices i and j , we obtain the linear map \mathcal{L} .

In this formulation, there are parameters S, K, T, L . For large enough T , we can indeed set $S = 1$ and still obtain accurate filter estimates as we will show below. The choice of the parameter K depends on the ensemble size that will be used in the actual online data assimilation. For the parameters T and L , one can choose L as large as possible depending on the computing capacity and T as large as possible depending on the length of the training data set or historical correlation statistics. In contrast with Anderson and Lei (2013), we do not need the historical true state to specify the map \mathcal{L} . In this sense, this data-driven method requires no tuning.

The complexity in constructing the correlation r_{xy_i} between a single scalar observation y_i and the state variable $x \in \mathbb{R}^N$ using an ensemble of size L is $\mathcal{O}(LN)$. Similarly, the cost for constructing the subsampled correlation using K members, where $K \leq L$, is $\mathcal{O}(KN)$. If these calculations are performed on historical data of length T with S -subsamples, then the complexity in constructing the regression matrix A is $\mathcal{O}(LNTS)$. In each minimization problem, the complexity in forming $A^\top A$ is $\mathcal{O}(N^2TS)$, in forming $A^\top b$ is $\mathcal{O}(NTS)$, and in inverting $A^\top A$ (say with LU factorization) is $\mathcal{O}(N^3)$. To avoid an underdetermined linear problem, $TS > N$, so $\mathcal{O}(LNTS)$ dominates, assuming that $L > N$. In our numerical experiment below, we will show accurate results even for $S = 1$. For high-dimensional problems one may consider a local least square fit, by assuming that each observation is only correlated to its nearby state space of dimension $N_{loc} < N$, as opposed to the whole state space. With these considerations, the complexity in solving each regression problem reduces to $\mathcal{O}(LN_{loc}T)$ and one can choose much smaller L to satisfy $N_{loc} < L < N$.

In the special case where $\mathcal{L}(\cdot, \cdot, j) \in \mathbb{R}^{N \times N}$ is restricted to be diagonal, the proposed minimization problem (5) seeks to estimate the localization function such that $\mathcal{L}(i, i, j)r^K(i, j) \approx r(i, j)$, where $\mathcal{L}(i, i, j) \in \mathbb{R}$; we refer to the resulting matrix from fitting to the diagonal \mathcal{L} as \mathcal{L}_d . In this case, the minimization procedure is pointwise and corresponds exactly to the strategies proposed in (Anderson 2012; Anderson and Lei 2013; Flowerdew 2015), with the only difference being in how the densities are estimated from the training data.

b. A serial EnKF with transformed correlation functions

In this paper, we will consider the implementation of the map \mathcal{L} on the serial local-least-squares LLS-EnKF (Anderson 2003), which assumes a diagonal observation error covariance R and uses the unperturbed observation component y_j^o to sequentially update the ensemble solutions. In each

assimilation step, the LLS-EnKF executes the following update

$$\begin{aligned}\hat{y}_j^a &= \hat{y}_j^f + (\hat{P}_{y_j y_j}^f + R_{jj})^{-1} \hat{P}_{y_j y_j}^f (y_j^o - \hat{y}_j^f), \\ y_j^{a,k} &= \hat{y}_j^a + \sqrt{\frac{R_{jj}}{R_{jj} + \hat{P}_{y_j y_j}^f}} (y_j^{f,k} - \hat{y}_j^f)\end{aligned}$$

and regresses the update $\Delta y_j^k = y_j^{a,k} - y_j^{f,k}$ onto the state variables as follows,

$$\mathbf{x}^{a,k} = \mathbf{x}^{f,k} + \frac{\hat{P}_{\mathbf{x} y_j}^f}{\hat{P}_{y_j y_j}^f} \Delta y_j^k. \quad (7)$$

The straightforward modification on the LLS-EnKF is to replace

$$\hat{P}_{\mathbf{x} y_j}^f = D_{X^f}^{1/2} r^K(\cdot, j) \sqrt{\hat{P}_{y_j y_j}^f} \longleftarrow D_{X^f}^{1/2} \mathcal{L}(\cdot, \cdot, j)^\top r^K(\cdot, j) \sqrt{\hat{P}_{y_j y_j}^f}, \quad (8)$$

in (7), where $D_{X^f} = \text{diag}(X^f X^{f\top}) / (K - 1)$ and $r^K(\cdot, j)$ are estimated from the EnKF with ensemble size K .

4. Numerical experiments on Lorenz-96 model

We test our methodology on an OSSE with the Lorenz-96 model (Lorenz 1996)

$$\frac{dx_j}{dt} = (x_{j+1} - x_{j-2})x_{j-1} - x_j + \mathcal{F}, \quad j = 0, \dots, N-1, \quad (9)$$

with $N = 40$ state variables ($\mathbf{x} \in \mathbb{R}^{40}$) equally spaced on a periodic one-dimensional domain of length 40 ($x_N = x_0$ and $x_{N+1} = x_1$). The forcing parameter $\mathcal{F} = 8$ is set for a strongly chaotic regime. A realization of the “true” signal or “nature run” \mathbf{x}^\dagger is obtained by integrating the system (9) forward in time, from random initial conditions, using a fourth-order Runge-Kutta scheme with a time step $\Delta t = 0.05$, or 6 hours.

Synthetic observations are generated from that truth trajectory based on the model

$$\mathbf{y}_m^o = h(\mathbf{x}_m^\dagger) + \boldsymbol{\epsilon}_m, \quad \boldsymbol{\epsilon}_m \sim \mathcal{N}(0, I), \quad (10)$$

at discrete time t_m , where ϵ_m , the observation noise, is unbiased and Gaussian with identity covariance matrix. Observations are made at uniform discrete times $t_{m+1} - t_m = n\Delta t$, where $n \in \mathbb{N}$ is the observation time step. In total 30 000 observations are collected, with the first 10 000 used for training purposes, and the last 20 000 for diagnostic evaluation. We will test both linear and non-linear observation operators $h: \mathbb{R}^N \rightarrow \mathbb{R}^M$, and in particular will focus on the following observation configurations:

- *Linear direct observations:* every p^{th} grid point is observed, where $p = N/M$, $M \in \{10, 20, 40\}$.
- *Linear indirect observations:* every other grid point is observed indirectly as a sum of itself and the 6 nearest neighbours. In this case, the linear observation operator H is defined such that

$$(H\mathbf{x})_j = \sum_{k=-3}^3 x_{\text{mod}(2j+k, N)}, \quad j = 1, \dots, M \quad (M = 20),$$

where the modulus is to preserve the periodicity of the domain.

- *Nonlinear indirect observations:* We consider a nonlinear observation function of the form

$$h_j(\mathbf{x}) = \sum_{k=-3}^3 w(x_{j+k})x_{j+k}, \quad j = 1, \dots, M \quad (M = 10), \quad (11)$$

with a state dependent weight, $w: [a, b] \rightarrow [0, 1]$ defined as

$$w(x_{j+k}) = \frac{c_{k+4}}{2} \left[1 + \cos \left(\frac{2\pi}{b-a} \left(x_j - \frac{a+b}{2} \right) \right) \right],$$

where $\mathbf{c} = (1, 0.8, 0.4, 0, 0.4, 0.8, 1)$ so when $k = \pm 3$, $c_{k+4} = 1$, etc. Here, the parameters a and b are the extrema of the data $\{\mathbf{x}_m\}_{m=1, \dots, 30000}$.

a. Localization map \mathcal{L}

Given some estimates r and r^K , the method in Section 3 provides an algorithm to learn the map \mathcal{L} offline with little computational overhead, by solving an appropriate minimization problem. Indeed for every pair (i, j) of state variable i and observation j , the map $\mathcal{L}(\cdot, i, j)$ is given as the solution of the least square problem (6), which is effectively a linear regression of the sample background correlation r^K onto r .

For the OSSE data assimilation system described by the model equations (9)-(10), the limiting correlation r needs to be estimated in some way. Here we use the output sample correlation of an ETKF with ensemble size 500 $(r_m^L = r_m^{500})_{m=1, \dots, T}$ at the first $T = 10000$ time steps as an estimator for $\{r_m\}_{m=1, \dots, T}$ (see a snapshot of r_m^L at an arbitrary instance m in Figure 1). Using the same data assimilation output, we also estimate the correlations $\{r_m^K\}_{m=1, \dots, T}$, subsampling K out of $L = 500$ members. In this numerical experiment, the ETKF is performed without any localization or inflation. Our choice of using ETKF is just to ensure that one can learn the map by fitting to correlation estimates from a different ensemble Kalman filtering method. In real applications, historical data or the output of a preferred cheap filter may be used as a proxy for the limiting correlation. The serial EnKF algorithm is then modified by applying the map $r^K \leftarrow \mathcal{L}^T r^K$ at each assimilation cycle, which transforms the undersampled correlation matrix r^K into a correlation matrix $\mathcal{L}^T r^K$ that is closer to r (see the corresponding snapshots in Figure 1).

Linear Observations. The maps learned in the training phase for linear direct and indirect observations are shown in Figures 2 and 3, respectively. The maps $\mathcal{L}(\cdot, i, j)$ are plotted for every pair (i, j) of state variables and observation at grid point $j = M/2$, where M is the number of observations ($M = 10, 40$). These curves are compared with the diagonal maps which arise from a separate minimization procedure, fitting to a diagonal $\mathcal{L}(\cdot, \cdot, j) \in \mathbb{R}^{N \times N}$; we refer to the resulting

map as \mathcal{L}_d . For direct observations, we note that both the map \mathcal{L} and the diagonal map \mathcal{L}_d peak to unity at the midpoint of the interval $[1, 40]$, which occurs precisely for the pair $(20, M/2)$ of state and observation variables, and slowly decay away from the observed location. We also find that the map \mathcal{L}_d can be viewed as an envelope for the packet of mapping curves \mathcal{L} . In this sense, the maps \mathcal{L} and \mathcal{L}_d capture the correlation structure between x and Hx , where H is the direct observation operator. The mapping curves are also more compact or “localized” when the ensemble size used is small, or when the observation time step is large. Interestingly, the regression produces maps with a structure that resembles the localization functions that are commonly used in EnKF applications; the key point is that here the linear map is not chosen a priori but instead is learned from the data. Based on these observations, we will refer to the maps \mathcal{L} and \mathcal{L}_d as *localization map* and *scalar localization function*, respectively.

Although the structure of the localization mappings can be easily anticipated in the case of direct observations, it may be more challenging to design a localization function when the state variables are observed indirectly. As shown in Figure 3, in the case of indirect observations, the localization maps exhibit a nontrivial structure that may be difficult to guess a priori. Interestingly, the maps have larger weights when the ensemble size is 40, compared to an ensemble of size 5, especially for smaller n . These weights also tend to decrease in magnitude when the observation time step increases. This suggests that when the ensemble size K is larger or when the observation time step n is small, the localization map gives more weights to the sample correlation r^K because it is a good estimator of r .

Nonlinear Indirect Observations. The maps for 10 nonlinear indirect observations observed at every 5 model steps ($n = 5$) are presented in Figure 4, for ensemble member of sizes 5 and 40. These maps show more complexity than the ones obtained in the linear indirect case, and reveal a particularly interesting trimodal envelope which is captured also in the diagonal maps \mathcal{L}_d . As

in the linear indirect case, the maps yield much greater weights when the ensemble size is 40, compared to when the ensemble is only 5.

Robustness with respect to the parameter S . It should be mentioned that, for the sake of computational efficiency, the localization maps shown in Figures 2, 3 and 4 were all obtained for $S = 1$ in the minimization problem (6), that is, when only one sample of the correlation r^K is taken from the sampling distribution $p(r^K|r)$. However, as outlined in Section 3, as many as $\binom{500}{K}$ samples can be obtained from randomly selecting K members from the ETKF analysis ensemble $\{x^{a,k}\}_{k=1}^{500}$. In Figure 5, we show the localization maps \mathcal{L} obtained from the minimization procedure with $S = 100$ correlation samples $\{r_s^K\}_{s=1}^{100}$, each one being computed by randomly selecting $K = 40$ analysis members from the 500 available, in the case of 10 linear direct observations taken at every model step. The resulting maps are arguably similar, at least qualitatively, to the ones previously shown in Figure 2 for the value $S = 1$, and the small residuals (pointwise difference) between the two sets of maps indicate that the method is not very sensitive to the parameter S .

b. Filtering results

We now verify the filtering performance, assimilating the last 20 000 observations, beyond the first 10 000 data set that is used for training. The performance of the filter is measured using the following metrics:

- The root mean squared error (RMSE) of the analysis ensemble mean from the truth x^\dagger , defined as

$$RMSE = \sqrt{\frac{1}{N} \sum_{j=1}^N (\hat{x}_j^a - x_j^\dagger)^2},$$

where \hat{x}_j^a and x_j^\dagger denote the j th component of the true state and the analysis mean estimate, respectively.

- The spread of the analysis from the ensemble mean, defined as

$$spread = \sqrt{\frac{1}{N(K-1)} \sum_{j=1}^N \sum_{k=1}^K (\delta x_j^{k,a})^2}$$

where $\delta x_j^{k,a} = x_j^{k,a} - \hat{x}_j^a$ is the deviation of the k^{th} analysis ensemble member from its ensemble mean.

Here we investigate the performance of the serial EnKF when the localization map \mathcal{L} is applied in conjunction with a multiplicative covariance inflation. The covariance inflation factor f is applied (Anderson 2001) to the prior ensemble estimate of the state covariance as follows $(1+f)^{P^b}$ in an attempt to reduce filter errors and avoid filter divergence.

Linear direct observations. Figure 6 shows the time mean of the RMSE (averaged over the last 20 000 cycles) in the case of linear direct observations, as a function of ensemble size and covariance inflation factor. The RMSE values reported here are normalized by the RMSE of the ETKF with 500 ensemble members, assimilating the same set of observations. Both the localization map \mathcal{L} and scalar localization function \mathcal{L}_d are tested on various number of observations ($M = 10, 20, 40$) and observation time steps ($n = 1, 5, 10$).

We see that compared to the benchmark ETKF solution, the serial EnKF filter using both the localization maps \mathcal{L} and \mathcal{L}_d performs well for large ensemble sizes, and its performance improves when the inflation factor is properly tuned. In the assimilation regime $M = 10$, $n = 10$, and $K = 5$, the filter's error at its lowest is about 17% more than the ETKF error, which is achieved when covariance inflation is completely switched off. Interestingly, if the observation time step is reduced to the model step ($n = 1$), the accuracy of the serial EnKF analysis' solution degrades to more than twice the ETKF error (top left colormap), even when the inflation factor is well tuned. In fact, in the regime $M = 10$ and $n = 1$, the filter diverges when the inflation factor goes beyond a threshold value of approximately 0.1. However, in many assimilation regimes (e.g. $M = 40$, $n = 1$,

and $K = 40$), the serial EnKF performs just as well as the benchmark ETKF. Notice also that in these direct observation cases, the sensitivity of the filter’s estimates to ensemble size and inflation factor using the map \mathcal{L} is not very different compared to that using the function \mathcal{L}_d (compare colormaps (a) and (b) in Figure 6).

In Figure 7, the time mean of RMSE for the best tuned inflation values are shown as a function of ensemble size, in the case of 10 and 20 direct linear observations, and observation time steps $n = 1$ and $n = 5$. In the same figure, we also show the corresponding time mean of spread. The figure compares three types of localization methods for the serial EnKF: the localization map \mathcal{L} , the scalar localization function \mathcal{L}_d , and the standard Gaspari-Cohn localization with an optimal half-width that is tuned in the training phase over a range of 1 to 10 grid points in each assimilation case (each case consisting of a specific choice of K , n , M , and inflation factor). We note that in most cases, the filter’s performance using the localization maps \mathcal{L} and \mathcal{L}_d is comparable with that using the GC localization function. In fact, when only 5 ensemble members are used, the localization maps \mathcal{L} and \mathcal{L}_d outperform GC in all instances shown except for the case $M = 20$, $n = 1$. In this later case, and for all larger ensemble sizes, GC still gives a more accurate filter solution. Nevertheless, these are still significant results, considering that the maps \mathcal{L} and \mathcal{L}_d are obtained offline without any tuning.

We found that the optimal GC full widths for the case of $n = 1, M = 10$ are 6, 16, 12, 16 for $K = 5, 10, 20, 40$, respectively. Comparing these optimal GC widths with the support of the localization maps in Figure 2 for the cases of $K = 5$ and $K = 40$, one can see that the map is negligible outside the support of the GC function.

In our cases of interest, i.e. for the localization maps \mathcal{L} and \mathcal{L}_d , the spread is close to the RMSE in most cases (but not always): the ensemble spread is large (more uncertainty about the estimator)

when the ensemble mean lies far from the truth, and the spread is small (more confidence about the estimator) when the ensemble mean lies close to the truth.

From the numerical results of Figure 7, notice that the serial EnKF with a well-tuned GC localization beats the “benchmark” ETKF in many cases when the ensemble size is large (closer to 40). At this point, one could ask whether a different choice of regressor for the localization maps \mathcal{L} and \mathcal{L}_d , such as the one obtained from the output of the serial EnKF with a well tuned GC localization, could produce an even more accurate filter solution. In Figure 8, we show the results using both the ETKF and the serial EnKF with GC localization as regressors, in the case of 20 linear direct observations taken at every model step. In the case of the GC regressor, the maps are trained on the correlation r obtained from 40 ensemble members of the serial EnKF with best tuned GC localization half-width parameter. The regressant r^K is then computed by selecting 5, 10, or 20 members out of these 40 ensemble members. Interestingly, the filter’s solution is more accurate using this new (GC) regressor, and the gain in accuracy is even more notable in the case of the scalar localization function \mathcal{L}_d . This is in fact a meaningful test, since it demonstrates how different regressors, such as the one obtained from the output of a cheap preferred filter, may be used without undermining the filtering skills of the localization map method.

Linear indirect observations. We next show in Figures 9 and 10 the results when 20 linear indirect observations are taken at observation time steps $n = 1, 5$ and 10. It is worth noting that when observations are taken at every model step ($n = 1$) and only 5 ensemble members are used, the performance of the map \mathcal{L} is remarkable compared to that of the GC localization: The RMSE for \mathcal{L} and GC are respectively 0.3602 and 5.0970, while the benchmark ETKF is at 0.1626 (see Figure 10). In this difficult regime, even the diagonal map \mathcal{L}_d outperforms GC with an RMSE of 3.3498. If one uses 10 members or more, all three localization methods perform qualitatively

just as well as ETKF. For instance, in the case of $K = 10$, the RMSE for \mathcal{L}_d , \mathcal{L} and GC are respectively 0.2033, 0.2182, and 0.2276, compared to 0.1626 for ETKF.

When observations are taken at every 5 model steps ($n = 5$) and the ensemble size is 5 or 10, the map \mathcal{L} offers a clear advantage over the other two methods. In fact when only 5 members are used, the filter with a GC localization produces a divergent solution, while the RMSE for the map \mathcal{L} is 2.2793, nearly four times the ETKF error (0.6369). In comparison, the diagonal map \mathcal{L}_d gives an error of 3.5728. However, for larger ensemble sizes, GC is outperforming the map \mathcal{L} , albeit not by much: for $K = 20$, the RMSE for GC, \mathcal{L} , and \mathcal{L}_d are respectively 0.4708, 0.7652, and 2.0077. For $K = 40$ the errors are, in the same order, 0.4145, 0.5722, and 0.7191.

When observations are taken at every $n = 10$ model steps, the localization maps \mathcal{L} and \mathcal{L}_d both outperform the GC localization (see Figure 10). In particular, the well-tuned GC estimates diverge with RMSEs above the climatological standard deviation (approximately 3.6), whereas the estimates from \mathcal{L} are much more accurate (or below the climatology). Also, as shown in Figure 9, the filter with either a localization map \mathcal{L} or \mathcal{L}_d produces a solution that is quite robust with respect to the inflation rate. In the case of the map \mathcal{L} and when 5 ensemble members are used, the error produced is about twice that obtained with the ETKF. That ratio goes down to about 1.65 when $K = 10$, 1.48 when $K = 20$, and 1.36 when $K = 40$. When the filter is used with the diagonal map \mathcal{L}_d , the error climbs to nearly 3 times the error of the benchmark ETKF in the case of 5 ensemble members.

Nonlinear indirect observations. In Figures 11 and 12, we compare the results for 10 nonlinear indirect observations as defined in (11) with observation time step $n = 5$. In this experiment, the localization map \mathcal{L} outperforms both \mathcal{L}_d and the well tuned GC localization for all ensemble sizes used. In Figure 11, we also include the sensitivity of the GC estimates with respect to inflation factor and ensemble size in the verification phase. As in all the experiments above, we should

mention again that the GC localization radius is obtained in the training phase. Notice that when only 5 ensemble members are used ($K = 5$), the filter’s solution with a GC localization diverges catastrophically (solution blows up), while the filter’s estimate gives an RMSE of 3.29 when a localization map \mathcal{L} is used. As the ensemble size increases, the filter estimates become closer to that of the benchmark ETKF. Even for $K = 10$, the GC filter’s estimate is very sensitive with respect to variance inflation.

5. Summary and conclusions

Ensemble Kalman filters provide the exact filter solution for linear models with Gaussian observational error in the limit of large ensemble sizes (Mandel et al. 2012). When the ensemble size is small, sampling errors tend to induce spurious correlations between observations and state variables. Localization methods address this issue by limiting the impact of an observation on a state variable, generally by specifying a radius of influence based on the physical distance between the two. However, most localizations require extensive tuning, which becomes impractical in large applications with multiscale spatial correlations (Dong et al. 2011; Kirchgessner et al. 2014; Flowerdew 2015). Moreover, when the observations are not correlated to the state variables by spatial distance, as may be the case in nonlinear models, the design of such localizations can be very difficult.

In this work we propose a data-driven localization method for the serial EnKF which assumes no functional dependence on physical distance or other imposed restrictions, and more importantly, it does not require any tuning. The localization is in fact “learned” from any time series of correlations – the regressor – that constitutes a good estimate of the filter limiting correlation, and is given as the solution of an appropriate linear regression problem. The localization takes indeed the form of a linear map \mathcal{L} that transforms, at each EnKF cycle, the poorly estimated sample correlation

into an improved correlation. We numerically tested the proposed methodology in an OSSE using the 40-variable Lorenz-96 model, with three different observation configurations: linear direct, linear indirect, and nonlinear indirect observations. We use as regressors the assimilation output of an ETKF with 500 ensemble members, and of a serial EnKF with a well-tuned GC localization with 40 ensemble members.

For linear direct observations, the linear maps \mathcal{L} and their diagonal counterparts \mathcal{L}_d have a well defined structure that captures the intrinsic correlations of the filter solution: they peak to unity at every observed state variable, and slowly decay away from that location. This structure is reminiscent of the localization functions that are commonly used in EnKF localization. For the indirect and nonlinear indirect observations, we obtain localization maps with bimodal and trimodal envelopes, respectively. From our experiments, we found that when the true correlation is local (as it is in the linear direct observation example) our method is not advantageous compared to the well-tuned GC. On the other hand, when the correlation is nonlocal (e.g. in the linear indirect and nonlinear indirect observation examples), our method consistently performs well, beating the well-tuned GC, since the localization map contains such nontrivial correlation information.

In the case of the ETKF regressor, we found that the linear map \mathcal{L} is more accurate and robust compared to both \mathcal{L}_d and the well-tuned GC localization method. In the case of the GC regressor, we found that the accuracy of the filter solutions with the localization maps \mathcal{L} and \mathcal{L}_d is improved compared to that obtained with the ETKF regressor. This test demonstrates that 1) the performance of the filter using the localization maps can be improved by using an even better correlation output as regressor and 2) the output of a cheap filter with small ensemble sizes can be used as a regressor when large ensemble data assimilation are too costly to be simulated, especially in large-dimensional applications. Therefore, the learning methodology proposed in this paper can be adopted using any available data assimilation correlation statistics and the results depend on the

quality of the training data, like in any data-driven method. Of course, the resulting map is also subject to noise but one can increase the number of data by fitting to the correlations with similar structure at different locations in the training phase. Finally, we should point out that one can use the nontrivial structure obtained from the training procedure here as a guideline to design a more appropriate parametric localization function.

While the localization map methodology was developed for the serial EnKF by correcting the correlations between observations and grid points based on learned data, we suspect that the same approach can be used on other variants of ensemble Kalman filters to improve the correlation statistical estimation. Also, the experiments reported in this paper are based on the perfect model assumption, while the proposed fitting strategy is not limited to OSSE. Thus it will be interesting to study the robustness of the proposed method in the presence of model errors. We plan to address these issues and implement this method on higher dimensional applications in our future research.

Acknowledgments. This research is partially supported by ONR MURI Grant No. N00014-12-1-0912, ONR Grant No. N00014-13-1-0797, and NSF Grant No. DMS-1317919.

References

- Anderson, J., 2001: An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, **129**, 2884–2903.
- Anderson, J., 2003: A local least squares framework for ensemble filtering. *Monthly Weather Review*, **131** (4), 634–642.
- Anderson, J., 2012: Localization and sampling error correction in ensemble Kalman filter data assimilation. *Monthly Weather Review*, **140** (7), 2359–2371.

- Anderson, J., and L. Lei, 2013: Empirical localization of observation impact in ensemble Kalman filters. *Monthly Weather Review*, **141** (11), 4140–4153.
- Anderson, J. L., 2016: Reducing correlation sampling error in ensemble Kalman filter data assimilation. *Monthly Weather Review*, **144**, 913–925.
- Bishop, C., B. Etherton, and S. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter part I: the theoretical aspects. *Monthly Weather Review*, **129**, 420–436.
- Bishop, C., and D. Hodyss, 2007: Flow-adaptive moderation of spurious ensemble correlations and its use in ensemble-based data assimilation. *Quarterly Journal of the Royal Meteorological Society*, **133** (629), 2029–2044.
- Casella, G., and R. Berger, 2002: *Statistical Inference*. Duxbury advanced series in statistics and decision sciences, Thomson Learning.
- Chen, Y., and D. Oliver, 2010: Cross-covariances and localization for enkf in multiphase flow data assimilation. *Computational Geosciences*, **14** (4), 579–601.
- Dong, J., M. Xue, and K. Droegemeier, 2011: The analysis and impact of simulated high-resolution surface observations in addition to radar data for convective storms with an ensemble kalman filter. *Meteorology and Atmospheric Physics*, **112** (1), 41–61.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, **99**, 10 143–10 162.
- Flowerdew, J., 2015: Towards a theory of optimal localisation. *Tellus A*, **67**.
- Furrer, R., and T. Bengtsson, 2007: Estimation of high-dimensional prior and posterior covariance matrices in kalman filter variants. *Journal of Multivariate Analysis*, **98** (2), 227 – 255.

- Gaspari, G., and S. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, **125**, 723–758.
- Hotelling, H., 1953: New light on the correlation coefficient and its transforms. *Journal of the Royal Statistical Society. Series B (Methodological)*, **15** (2), 193–232.
- Houtekamer, P., and H. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Monthly Weather Review*, **126**, 796–811.
- Hunt, B., E. Kostelich, and I. Szunyogh, 2007: Efficient data assimilation for spatiotemporal chaos: a local ensemble transform Kalman filter. *Physica D*, **230**, 112–126.
- Jun, M., I. Szunyogh, M. G. Genton, F. Zhang, and C. H. Bishop, 2011: A statistical investigation of the sensitivity of ensemble-based kalman filters to covariance filtering. *Monthly Weather Review*, **139** (9), 3036–3051.
- Kirchgeßner, P., L. Nerger, and A. Bunse-Gerstner, 2014: On the choice of an optimal localization radius in ensemble kalman filter methods. *Monthly Weather Review*, **142** (6), 2165–2175.
- Lorenc, A., 2003: The potential of the ensemble Kalman filter for NWP—a comparison with 4D-Var. *Quarterly Journal of Royal Meteorological Society*, **129**, 3183–3203.
- Lorenz, E. N., 1996: Predictability: A problem partly solved. *Proc. Seminar on predictability*, Vol. 1.
- Mandel, J., L. Cobb, and J. D. Beezley, 2012: On the convergence of the ensemble Kalman filter. *Applications of Mathematics*, **56** (6), 533–541.
- Zhen, Y., and F. Zhang, 2014: A probabilistic approach to adaptive covariance localization for serial ensemble square root filters. *Monthly Weather Review*, **142** (12), 4499–4518.

LIST OF FIGURES

- Fig. 1.** The map \mathcal{L} transforms the undersampled background error correlation matrix r^K into a sample of correlation matrix $\mathcal{L}^T r^K$ resembling more the regressor r^L . A snapshot of all three correlation matrices are shown here for 20 linear direct observations taken every model time step, for an ensemble of size $K = 5$. The regression is based on the output sample correlation of an ETKF with ensemble size $L=500$ 28
- Fig. 2.** The map $\mathcal{L}(q, i, j)$ for the middle observation $j = M/2$, shown as a function of the state variables $q = 1, \dots, 40$ for $M = 10$ and 40 linear direct observations and for observation time steps $n = 1, 10$. The map \mathcal{L} is plotted here for every pair (i, j) of state variable $i = 1, \dots, 40$ and the middle observation $j = M/2$ in the case of (a) 5 and (b) 40 ensemble members. The corresponding diagonal map \mathcal{L}_d is overlaid in dashed black. 29
- Fig. 3.** Same as in Figure 2 but for $M = 20$ indirect observations and observation time steps $n = 1, 5$, and 10. 30
- Fig. 4.** Same as in Figure 2, but for $M = 10$ nonlinear indirect observations and observation time step $n = 5$ 31
- Fig. 5.** (Left) Maps \mathcal{L} obtained using the value $S = 100$, in the case of 10 linear direct observations taken at every model step, and for ensemble size 40. These maps should be compared with those of Figure 2b ($M = 10$ and $n = 1$) obtained with $S = 1$. (Right) Residuals (pointwise difference) between the maps obtained with the values $S = 1$ and $S = 100$ 32
- Fig. 6.** *Linear direct observations.* Array of time mean RMSE colormaps (normalized by the RMSE of the ETKF using 500 ensemble members) for the serial EnKF using (a) the localization map \mathcal{L} and (b) the scalar localization function \mathcal{L}_d , plotted against inflation factor and ensemble size $K = 5, 10, 20, 40$. The 3 by 3 array's columns and rows correspond to observation time step n and number M of linear direct observations, respectively. White pixels correspond to RMSE values outside the scale shown. 33
- Fig. 7.** *Linear direct observations.* Time mean of RMSE and spread for the serial EnKF as a function of ensemble size using three types of localization: the localization map \mathcal{L} , the scalar localization function \mathcal{L}_d , and the Gaspari-Cohn (GC) localization with an optimal half-width. All use the best tuned inflation values. These are compared with the benchmark ETKF solution using 500 ensemble members. Also shown is the observation error variance. 34
- Fig. 8.** Time mean of RMSE and spread as a function of ensemble size for the serial EnKF with localization maps \mathcal{L} and \mathcal{L}_d , obtained from two different choices of regressors: 1) the ETKF with 500 ensemble members, and 2) the serial EnKF with GC localization with optimal half-width. All use the best tuned inflation values. The results are for 20 linear direct observations taken at every model step. 35
- Fig. 9.** *Linear indirect observations.* Same as in Figure 6, but for 20 linear indirect observations and with a different scale. 36
- Fig. 10.** *Linear indirect observations.* Same as in Figure 7, but for 20 linear indirect observations, and observation time steps $n = 1, 5$, and 10. 37
- Fig. 11.** Time mean RMSE for the serial EnKF in the case of 10 nonlinear indirect observations and observation time step $n = 5$, using three different localization methods: the localization maps \mathcal{L} and \mathcal{L}_d , and the Gaspari-Cohn (GC) localization function with optimal half-width.

The RMSE is normalized by the RMSE of the ETKF using 500 ensemble members, and plotted against inflation factor and ensemble size. White pixels correspond to RMSE values outside the scale shown. 38

Fig. 12. *Nonlinear Indirect Observations.* Same as in Figure 7, but for 10 nonlinear indirect observations and observation time step $n = 5$ 39



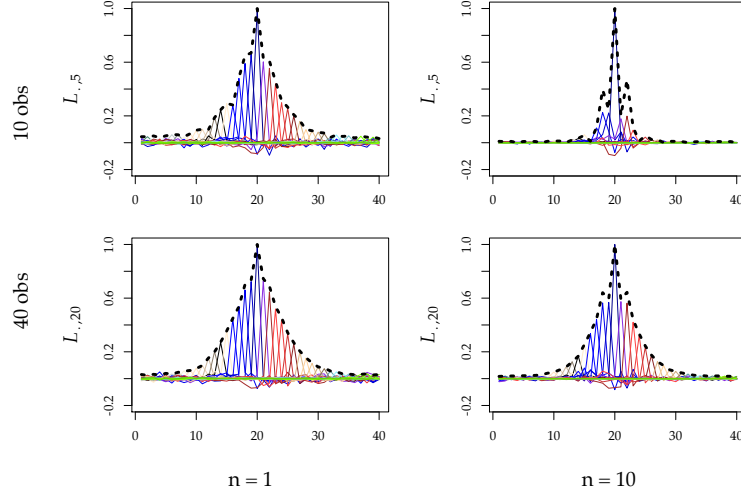


FIG. 1. The map \mathcal{L} transforms the undersampled background error correlation matrix r^K into a sample of correlation matrix $\mathcal{L}^T r^K$ resembling more the regressor r^L . A snapshot of all three correlation matrices are shown here for 20 linear direct observations taken every model time step, for an ensemble of size $K = 5$. The regression is based on the output sample correlation of an ETKF with ensemble size $L=500$.




(a) 5 ensemble members.

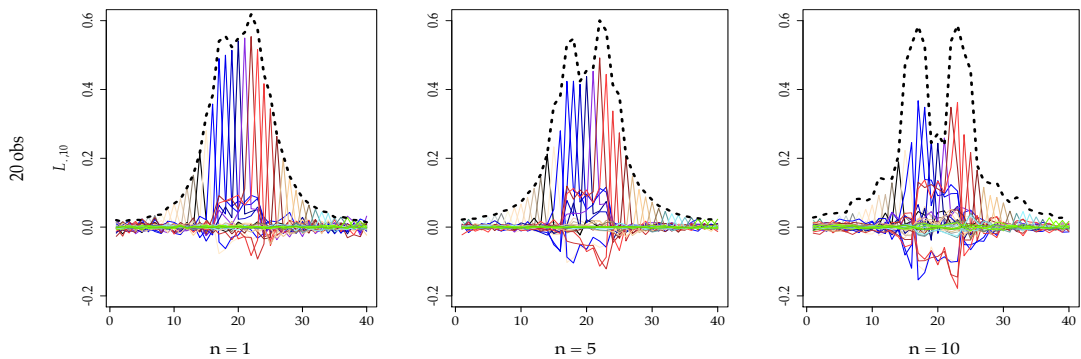


(b) 40 ensemble members.

FIG. 2. The map $\mathcal{L}(q, i, j)$ for the middle observation $j = M/2$, shown as a function of the state variables $q = 1, \dots, 40$ for $M = 10$ and 40 linear direct observations and for observation time steps $n = 1, 10$. The map \mathcal{L} is plotted here for every pair (i, j) of state variable $i = 1, \dots, 40$ and the middle observation $j = M/2$ in the case of (a) 5 and (b) 40 ensemble members. The corresponding diagonal map \mathcal{L}_d is overlaid in dashed black.



(a) 5 ensemble members.



(b) 40 ensemble members.

FIG. 3. Same as in Figure 2 but for $M = 20$ indirect observations and observation time steps $n = 1, 5$, and 10 .




FIG. 4. Same as in Figure 2, but for $M = 10$ nonlinear indirect observations and observation time step $n = 5$.






FIG. 5. (*Left*) Maps \mathcal{L} obtained using the value $S = 100$, in the case of 10 linear direct observations taken at every model step, and for ensemble size 40. These maps should be compared with those of Figure 2b ($M = 10$ and $n = 1$) obtained with $S = 1$. (*Right*) Residuals (pointwise difference) between the maps obtained with the values $S = 1$ and $S = 100$.



(a) Localization map \mathcal{L} .



(b) Scalar localization function \mathcal{L}_d .

FIG. 6. *Linear direct observations.* Array of time mean RMSE colormaps (normalized by the RMSE of the ETKF using 500 ensemble members) for the serial EnKF using (a) the localization map \mathcal{L} and (b) the scalar localization function \mathcal{L}_d , plotted against inflation factor and ensemble size $K = 5, 10, 20, 40$. The 3 by 3 array's columns and rows correspond to observation time step n and number M of linear direct observations, respectively. White pixels correspond to RMSE values outside the scale shown.




FIG. 7. *Linear direct observations.* Time mean of RMSE and spread for the serial EnKF as a function of ensemble size using three types of localization: the localization map \mathcal{L} , the scalar localization function \mathcal{L}_d , and the Gaspari-Cohn (GC) localization with an optimal half-width. All use the best tuned inflation values. These are compared with the benchmark ETKF solution using 500 ensemble members. Also shown is the observation error variance.






FIG. 8. Time mean of RMSE and spread as a function of ensemble size for the serial EnKF with localization maps \mathcal{L} and \mathcal{L}_d , obtained from two different choices of regressors: 1) the ETKF with 500 ensemble members, and 2) the serial EnKF with GC localization with optimal half-width. All use the best tuned inflation values. The results are for 20 linear direct observations taken at every model step.



(a) Localization map \mathcal{L} .



(b) Scalar localization function \mathcal{L}_d .

FIG. 9. *Linear indirect observations*. Same as in Figure 6, but for 20 linear indirect observations and with a different scale.

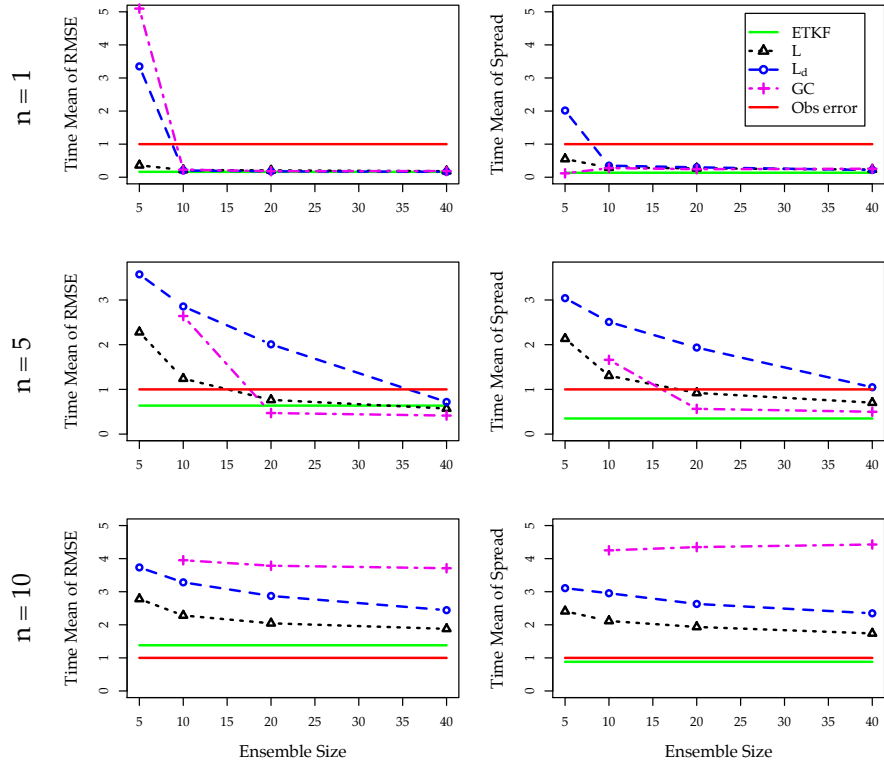


FIG. 10. *Linear indirect observations.* Same as in Figure 7, but for 20 linear indirect observations, and observation time steps $n = 1, 5$, and 10 .




FIG. 11. Time mean RMSE for the serial EnKF in the case of 10 nonlinear indirect observations and observation time step $n = 5$, using three different localization methods: the localization maps \mathcal{L} and \mathcal{L}_d , and the Gaspari-Cohn (GC) localization function with optimal half-width. The RMSE is normalized by the RMSE of the ETKF using 500 ensemble members, and plotted against inflation factor and ensemble size. White pixels correspond to RMSE values outside the scale shown.




FIG. 12. *Nonlinear Indirect Observations*. Same as in Figure 7, but for 10 nonlinear indirect observations and observation time step $n = 5$.